

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

THIS PAGE BLANK (USPTO)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
2 May 2002 (02.05.2002)

PCT

(10) International Publication Number
WO 02/35345 A2

(51) International Patent Classification⁷: G06F 9/00

(21) International Application Number: PCT/US01/51236

(22) International Filing Date: 25 October 2001 (25.10.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/697,222 26 October 2000 (26.10.2000) US

(71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901
San Antonio Road, M/S PALO1-521, Palo Alto, CA 94303
(US).

(72) Inventors: NGUYEN, Hien, H.; 282 Grove Street, Apt. 1,
Auburndale, MA 02466 (US). MORRIER, Donald, M.;
PMB 312, 119 Drum Hill Road, Chelmsford, MA 01824
(US).

(74) Agents: ASHER, Robert, M. et al.; Bromberg & Sunstein
LLP, 125 Summer Street, Boston, MA 02110-1618 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM,
HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK,
LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX,
MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK,
SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

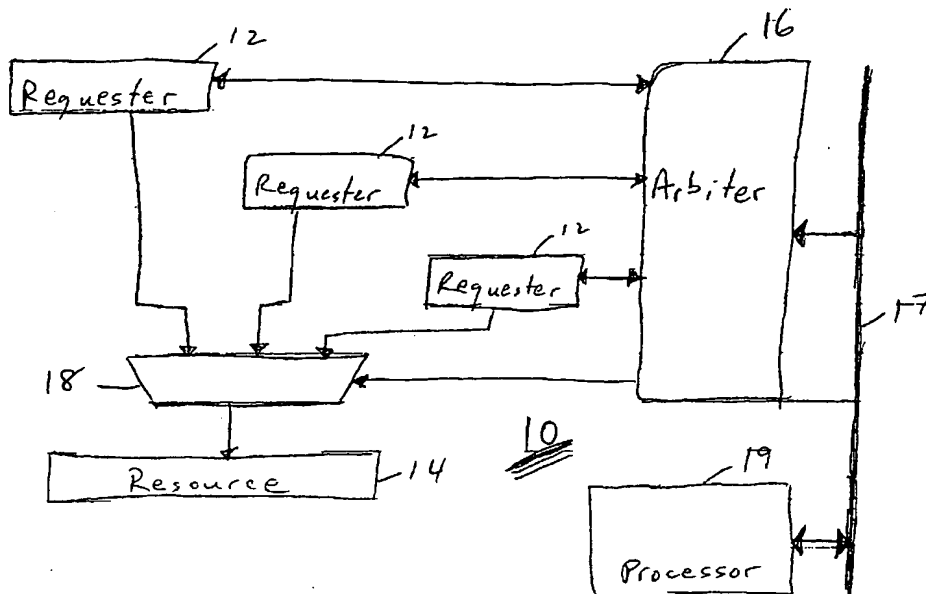
(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian
patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European
patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,
IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF,
CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD,
TG).

Published:

— without international search report and to be republished
upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: ARBITRATION METHOD AND APPARATUS



(57) Abstract: An arbitration method and mechanism assigning priority to one of a combination of requesters. A priority vector for the combination of requesters has a sequential list of requester identifiers. Upon receiving requests from the combination of requesters, the corresponding priority vector is referred to and the next requester identifier on this list is selected. Priority is awarded to the requester corresponding to that requester identifier. A priority vector is preferably provided for each possible combination of multiple requesters. The list of requester identifiers in each priority vector may be advantageously programmable. Preferably, all requester identifiers for a single requester are not provided in a contiguous clump within a given priority vector.

WO 02/35345 A2

AM

Background of the Invention

The present invention relates to an arbitration mechanism for awarding priority to one of a plurality of requesters, in particular, such a mechanism which guarantees a minimum software programmable slice of available bandwidth to each requestor.

- 5 Where there is a single resource in the data processing system whether that resource be a bus, an SRAM, a FIFO, etc., an arbitrator is required to control access to that resource. In accordance with many arbitration schemes, each requestor has an assigned priority. Requesters with the same priority are given access in an equitable or round-robin fashion. Typically, requests with higher priority are given preferred access,
- 10 preempting requests with lower priority. Additional mechanisms have been developed, that when added on top of such a system, prevent requesters with low priority from being starved from access to the resource.

Summary of the Invention

- 15 In accordance with an arbitration method embodiment of the invention, requests for access to a resource are received from a combination of requesters. A priority vector associated with that combination of requesters is referred to in order to identify the requestor to be awarded priority. Requestor identifiers each corresponding to one of the requesters are listed sequentially in the priority vector. The number of requestor
- 20 identifiers in the vector exceed the number of requesters in the combination. A pointer identifies a next requester identifier on the list. The next requester identifier is selected and the requester corresponding to that requester identifier is awarded priority. The pointer may be incremented sequentially through the list of requester identifiers. Programmability can be made available for the total number of requester identifiers in
- 25 the priority vector and for the number of requester identifiers corresponding to each requester.

According to a further embodiment of the arbitration method, a plurality of priority vectors are provided, one for each possible combination of multiple requesters that may seek access to the resource. The priority vector is a sequence of requester

identifiers, each requester identifier corresponding to one of the requesters. When requests for access are received from the combination of requesters, the priority vector corresponding to the given combination of requesters is used for selection of the next requester identifier. Priority is awarded to the requester corresponding to the next

5 requester identifier.

In accordance with embodiments of an arbitration mechanism of the invention, a selector, responsive to requests for access from the requesters in a given combination of requesters, ascertains a next requester identifier from the sequence of requester identifiers in the priority vector corresponding to the given combination of requesters.

10 Further, a pointer may be included for pointing to one of the requester identifiers in a priority vector. The sequence of requester identifiers may be held in a register.

A data processing system embodiment of the invention includes a resource, a plurality of requesters and an arbiter in communication with the requesters. The arbiter includes a selector responsive to the requests for ascertaining a next requester identifier
15 from an appropriate priority vector and awarding priority to the requester corresponding to a next requester identifier.

Since each requester will have at least one requester identifier in the priority vector used to award priority, even low priority requesters, will be able to gain some minimal access to the resource. The ability to assign a percentage of bandwidth to each
20 requester by adjusting the number of corresponding requester identifiers in priority vectors is an advantage of various embodiments of the invention.

Other objects and advantages of the invention will become apparent during the following description of the presently preferred embodiments of the invention taken in conjunction with the drawings.

25

Brief Description of the Drawings

FIG. 1 is a block diagram of a data processing system of an embodiment of the present invention.

FIG. 2 is a schematic block diagram illustrating an arbitration method and

mechanism of embodiments of the present invention.

Detailed Description of the Preferred Embodiments

Referring now to FIG. 1, in the data processing system 10 it is often the case that
5 a plurality of requesters 12 may at the same time request access to a single resource 14. The resource may be any of a variety of components such as a bus, an SRAM, a FIFO, a network communication port, a graphic processing pipeline, DRAM memory banks, a data structure in a data base program, etc. In order to decide which requester receives access to the single resource 14, an arbiter 16 is provided.

10 Arbiter 16 receives requests for access from any of a number of requesters. The arbiter 16 awards priority to one of the requesters. In a preferred embodiment, arbiter 16 sends a select signal to multiplexor 18 which selects the data from the requester awarded priority. The selected data proceeds to the resource 14. Also in this embodiment, arbiter 16 informs the requester awarded priority that its request was granted. Preferably, the
15 arbiter is programmable. A processor 19 is shown for configuring the arbiter through a bus 17.

A preferred embodiment of an arbiter of the present invention is illustrated in FIG. 2. The arbitration method is required when more than one requester 12 seeks access to a single resource at one time. For any given combination of requesters, the
20 arbiter 16 has a priority vector 21 with a sequential list of requester identifiers. A requester identifier corresponds to one of the requesters. In a preferred embodiment, the number of requester identifiers in the list of the priority vector 21 exceeds the number of requesters in the given combination. In accordance with further embodiments, each possible combination of multiple requesters is associated with a priority vector 21. For
25 each priority vector 21, a pointer 24 or equivalent designator identifies the next requester identifier on the list. For the given combination of multiple requesters, the associated priority vector is referred to and the next requester identifier from that priority vector is selected. Priority is awarded to the requester corresponding to the next requester identifier. The pointer 24 is then changed to point to the next requester identifier along

the sequence of requester identifiers. As described, the arbitration method may be performed entirely in software or in a combination of software and hardware as described below with respect to FIG. 2.

The arbitration mechanism includes a selector responsive to requests for access.

- 5 The selector is responsible for ascertaining the next requester identifier in the appropriate priority vector for awarding priority to the requester corresponding to the next requester identifier. The selector in FIG. 2 includes a multiplexor 20 for selecting the appropriate priority vector and a multiplexor 22 for selecting the appropriate pointer. The multiplexor 20 receives an input for each possible requester indicating whether a request
- 10 is pending for the single resource. The embodiment shown is for a system with a maximum of three requesters. Therefore, there is a priority vector 21 for all three requesters and three additional priority vectors 21 for each possible combination of two requesters. Based on which requesters have requests pending, the multiplexor 20 immediately selects the corresponding priority vector 21. In the example of the
- 15 illustration, the first and second requesters have pending requests and the priority vector for requester 1 and requester 2 is selected.

- Priority is awarded to the requester corresponding to the next requester identifier in the priority vector. The priority vector is a sequential list of requester identifiers, each requester identifier corresponding to one of the requesters. In the illustration, the
- 20 requester identifier is simply the number 1, 2 or 3 corresponding respectively to requester 1, requester 2 or requester 3. In the priority vector for requester 1 and requester 2, the list includes requester identifiers corresponding to the two requesters 1 and 2. In the illustration, the priority vector for requester 1 and 2 split the priority 50/50 between requester 1 and requester 2, each one having five requester identifiers in the list. In the
- 25 case where a requester has more than one requester identifier in the list, it is desirable not to clump all of its requester identifiers together as such arrangement may inordinately starve the other requester(s) out over a clump of time. Instead, it is preferred for each requester having more than one requester identifier, that its requester identifiers not be contiguous all to one another. This distributes the grant of priority to the various

requesters more evenly over time.

The next requester identifier in a priority vector is indicated by a pointer 24. The multiplexor 22 receives the inputs indicating which requesters have a pending request and uses those inputs to immediately select the pointer corresponding to the selected
5 priority vector. The pointer 24 is used to ascertain the next requester identifier which in turn determines which requester is awarded priority. The award of priority is signaled back to the requester so it knows its request is granted. The award of priority is also used by the resource to select the winning request and process it. In FIG. 1, this is through multiplexor 18. Upon using a pointer 24 to select the award of priority, the
10 pointer needs to be shifted to point to the following requester identifier on the priority vector. This may be accomplished by simply incrementing the pointer. A vector length indicator (numslot) 26 accompanies the pointer for each priority vector. When the total number of requester identifiers in the priority vector have been traversed by the pointer as determined by the vector length indicator, the pointer will wrap around to the first
15 requester identifier on the list. A pointer control module 28 receives the pointer and associated vector length indicator 26 so as to generate the next pointer value and substitute it into the pointer 24. The pointer control module 28 can be formed in hardware, software or anything in between.

In accordance with an embodiment of the present invention, the number of
20 requester identifiers in the priority vector exceeds the number of requesters. Thus, for a system with three possible requesters, the priority vector for all three requesters exceeds three requester identifiers. The more requester identifiers in a priority vector, the smaller the minimum bandwidth which can be assigned to a low priority requester. A priority vector with 10 requester identifiers can assign as little as one requester identifier to a
25 requester giving it approximately a 10% bandwidth. A priority vector with 16 requester identifiers permits assigning a bandwidth as small as about 6% to a low priority requester.

In accordance with an embodiment of the invention, the number of requester identifiers in any given priority vector is programmable. Software instructions can be

used to set the vector length for each of the priority vectors. Further programmability is also available in that the proportion of requester identifiers assigned to the requesters can be varied for each priority vector. Although each priority vector can be individually programmed with a desired sequence and proportion of requester identifiers, it is

5 generally preferable to determine a desired percentage bandwidth for each requester and have the various priority vectors attempt to achieve that goal. A system that desires to provide about 40% access to requester 1, about 40% access to requester 2 and about 20% access to requester 3 with a vector length of 10 for each priority vector is shown in FIG.

2. The priority vector for all three provides four requester identifiers for requester 1,
10 four requester identifiers for requester 2 and two requester identifiers for requester 3. The priority vector for requesters 1 and 2 correspondingly provide five requester identifiers for requester 1 and five requester identifiers for requester 2. The priority vector for requesters 1 and 3 provides six requester identifiers for requester 1 and four requester identifiers for requester 3. Similarly the priority vector for requesters 2 and 3
15 provide six requester identifiers for requester 2 and four requester identifiers for requester 3. One might also have chosen to divide these vectors seven to three rather than six to four. As there are a small finite number of requester identifiers, bandwidth is only proportioned approximately. Increasing vector length can make the proportions more precise. Although it is preferred to have consistency of proportions among the
20 priority vectors, it is within the scope of the invention that each priority vector be independently programmed with its own list of requester identifiers without regard to overall proportions. Each requester can thus be given a desired number of requester identifiers in the priority vectors. In each priority vector, the requester identifiers for any given requester are preferably spread across the vector and do not appear in a single
25 contiguous clump.

The arbiter 16 of a presently preferred embodiment can be implemented on an ASIC. The priority vectors may be stored in a register. Each register may be divided into fields. One field provides slots that are filled by the requester identifiers to form priority vectors. A second field indicates the vector length, i.e., the number of slots that

are filled with requester identifiers. A large register can hold several priority vectors in separate fields and have associated fields which indicate the vector lengths for the corresponding priority vectors. The contents of the registers are programmable and may be filled to satisfy the requirements of a user through a processor 19 for example.

- 5 Although shown and described with respect to hardware such as the multiplexors and registers, embodiments of the invention may also be made implementing aspects of the invention in software.

- Although the embodiments hereinbefore described are preferred, many modifications and refinements which do not depart from the true spirit and scope of the invention may be conceived by those skilled in the art. For example, the requester identifiers need not be numbers that match their corresponding requester but may be assigned via any desired coding. Shifting the pointer along a priority vector may involve any of incrementing, decrementing or other algorithm for moving the pointer in sequence through a list of requester identifiers. It is intended that all such modifications, including
- 10 but not limited to those set forth above, be covered by the following claims.
- 15

WE CLAIM:

1. A method of arbitrating requests for access to a resource for use in a data processing system having a plurality of requesters and the resource, said method comprising:

- 5 receiving requests for access from a combination of said requesters;
referring to a priority vector associated with said combination of requesters, each priority vector having a sequential list of requester identifiers, the number of requester identifiers in the list exceeding the number of requesters in said combination, and a pointer to a next requester identifier on the list;
- 10 selecting the next requester identifier from the priority vector associated with said combination of requesters; and
awarding priority to the requester corresponding to the next requester identifier.

2. The method of claim 1 further comprising setting the number of requester
15 identifiers in the priority vector.

3. The method of claim 2 further comprising setting how many of the requester identifiers correspond to each requester in the combination of requesters associated with the priority vector.

20 4. The method of claim 3 wherein for those requesters having a plurality of requester identifiers in the priority vector, the corresponding requester identifiers for each of those requesters not all being contiguous in the sequential list.

25 5. The method of claim 1 wherein said act of referring comprises identifying from among a plurality of priority vectors, the priority vector associated with said combination of requesters.

6. The method of claim 1 comprising sequentially changing the pointer along the sequence of requester identifiers each time a next requester identifier is selected.

7. A method of arbitrating requests for access to a resource for use in a data
5 processing system having a plurality of requesters and the resource, said method comprising:

providing a plurality of priority vectors, one priority vector for each possible combination of multiple requesters seeking access to the resource, each vector being a sequence of requester identifiers, each requester identifier corresponding to one of the
10 requesters;

receiving requests for access to the resource from a given combination of said requesters;

selecting a next requester identifier from the priority vector corresponding to the given combination of requesters;

15 awarding priority to the requester corresponding to the next requester identifier.

8. The method of claim 7 further comprising assigning each requester in the combination of requesters corresponding to one of the priority vectors a chosen number of requester identifiers in the priority vector.

20

9. The method of claim 8 wherein for those requesters having a chosen number greater than one for one of said priority vectors, the corresponding requester identifiers for each of those requesters not all being contiguous in the sequence of requester identifiers.

25

10. The method of claim 7 further comprising providing a pointer for each priority vector that points to the next requester identifier in the priority vector and sequentially changes along the sequence of requester identifiers each time a next requester identifier is selected.

11. A method for providing at least minimal access to a resource for each requester among a plurality of requesters, said method comprising:

5 providing a plurality of priority vectors, each priority vector corresponding to a possible combination of said requesters and being capable of holding a set number of requester identifiers in a sequence;

programming each priority vector such that it contains a desired number of requester identifiers corresponding to each of the requesters in the corresponding combination of requesters;

10 pointing to a next requester identifier in each of said priority vectors;

awarding priority to a requester from a combination of said requesters by identifying the next requester in the priority vector corresponding to the combination of requesters and pointing to a requester identifier following the next requester identifier in the sequence of said priority vector.

15

12. The method of claim 11 wherein for each requester having a number of requester identifiers greater than one in one of said priority vectors, the corresponding requester identifiers not all being contiguous in the sequence of requester identifiers of said one of said priority vectors.

20

13. The method of claim 11 further comprising programming the set number of requester identifiers for each priority vector.

14. The method of claim 11 wherein pointing to a requester identifier following the next requester identifier is performed in a wrap around manner upon reaching the end of the sequence.

25

15. An arbitration mechanism for arbitrating requests for access to a resource comprising:

a plurality of priority vectors, each providing a sequence of requester identifiers corresponding to requesters, the requester identifiers in each priority vector corresponding to requesters forming a different combination of requesters;

- a selector, responsive to requests for access from the requesters in a given
5 combination of requesters, for ascertaining a next requester identifier in sequence in the priority vector corresponding to the given combination of requesters and awarding priority to the requester corresponding to the next requester identifier.

16. The arbitration mechanism of claim 15 wherein each priority vector has a number
10 of requester identifiers exceeding the number of requesters in the respective combination of requesters.

17. The arbitration mechanism of claim 15 wherein for each requester having a
number of requester identifiers greater than one in one of said priority vectors, the
15 corresponding requester identifiers not all being contiguous in the sequence of requester identifiers of said one of said priority vectors.

18. The arbitration mechanism of claim 15 further comprising a processor that sets
the total number of requester identifiers in each priority vector and sets the number of
20 requester identifiers corresponding to any given requester in each priority vector.

19. An arbitration mechanism comprising:
a register for holding a sequence of requester identifiers, said register including
requester identifiers corresponding to each requestor of a combination of requesters,
25 wherein the number of requester identifiers in said register exceeds the number of requesters in said combination;
a pointer for pointing to one of the requester identifiers in said register; and
a selector, responsive to requests from the requesters in the combination of requesters, for awarding priority to the requester corresponding to the requester identifier

being pointed at by said pointer; and

means for shifting said pointer along the sequence of requester identifiers.

20. The arbitration mechanism of claim 19 wherein said pointer shifts sequentially
5 through said sequence of requester identifiers in a wrap around manner.

21. The arbitration mechanism of claim 19 further comprising a plurality of said
registers, each including requester identifiers corresponding to a different combination of
requesters.
10

22. The arbitration mechanism of claim 21 wherein said selector responds to
requests from a given combination of requesters by awarding priority using the register
having requester identifiers corresponding to the given combination of requesters.

15 23. A data processing system comprising:
a resource;
a plurality of requesters;
an arbiter in communication with said plurality of requesters, said arbiter
including:

20 a plurality of priority vectors, each providing a sequence of requester
identifiers, each requester identifier corresponding to one of said requesters, the
requester identifiers in each priority vector corresponding to requesters from a
different combination of requesters;

25 a selector, responsive to requests for access from the requesters in
a given combination of said requesters, for ascertaining a next requester
identifier in sequence in the priority vector corresponding to the given
combination of requesters and awarding priority to the requester
corresponding to the next requester identifier.

30 24. The data processing system of claim 23 wherein each priority vector has a

number of requester identifiers exceeding the number of requesters in the respective combination of requesters.

25. The data processing system of claim 23 wherein for each requester having a
5 number of requester identifiers greater than one in one of said priority vectors, the corresponding requester identifiers not all being contiguous in the sequence of requester identifiers of said one of said priority vectors.

26. The data processing system of claim 23 further comprising a processor that sets
10 the total number of requester identifiers in each priority vector and sets the number of requester identifiers corresponding to any given requester in each priority vector.

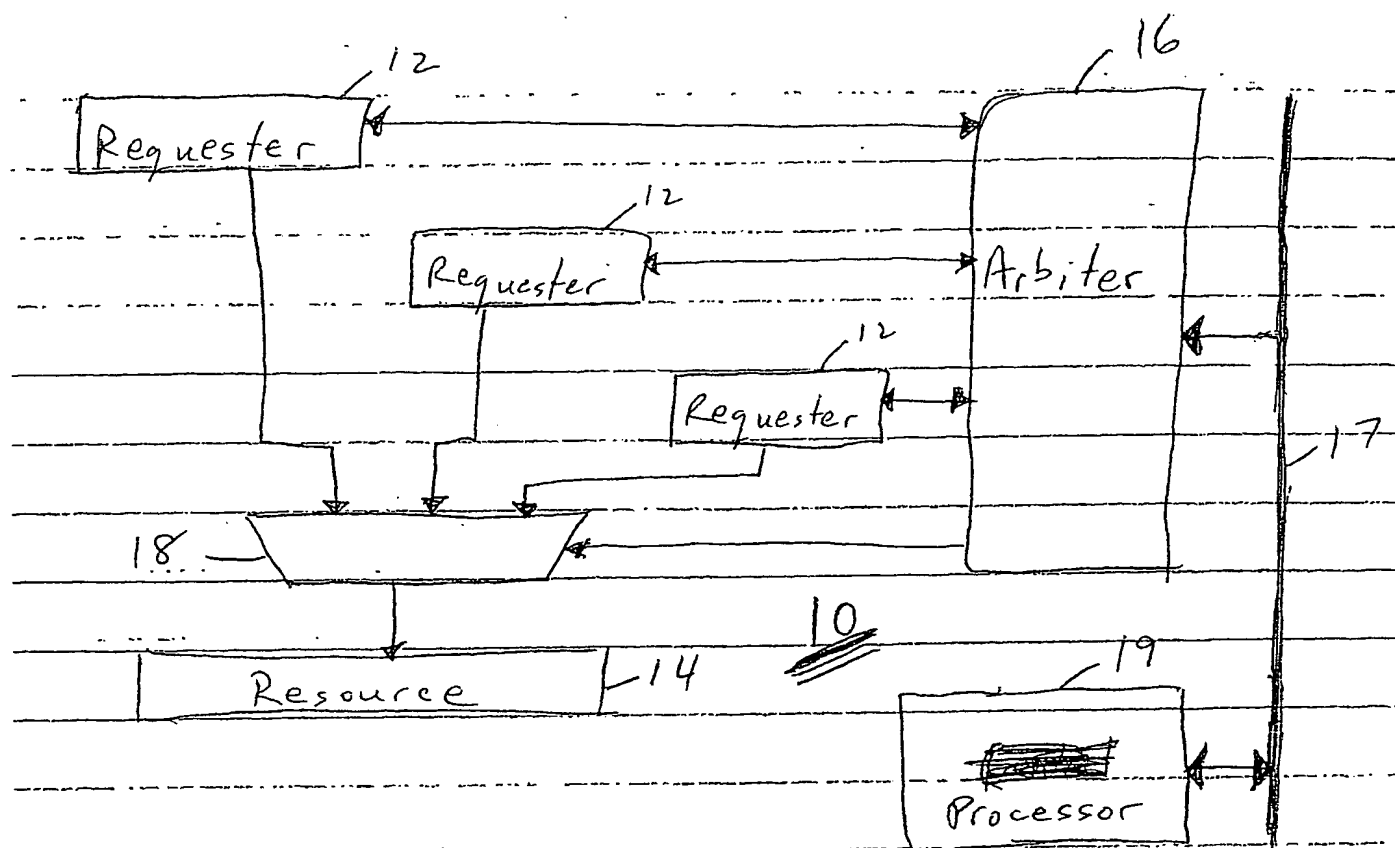


FIG. 1

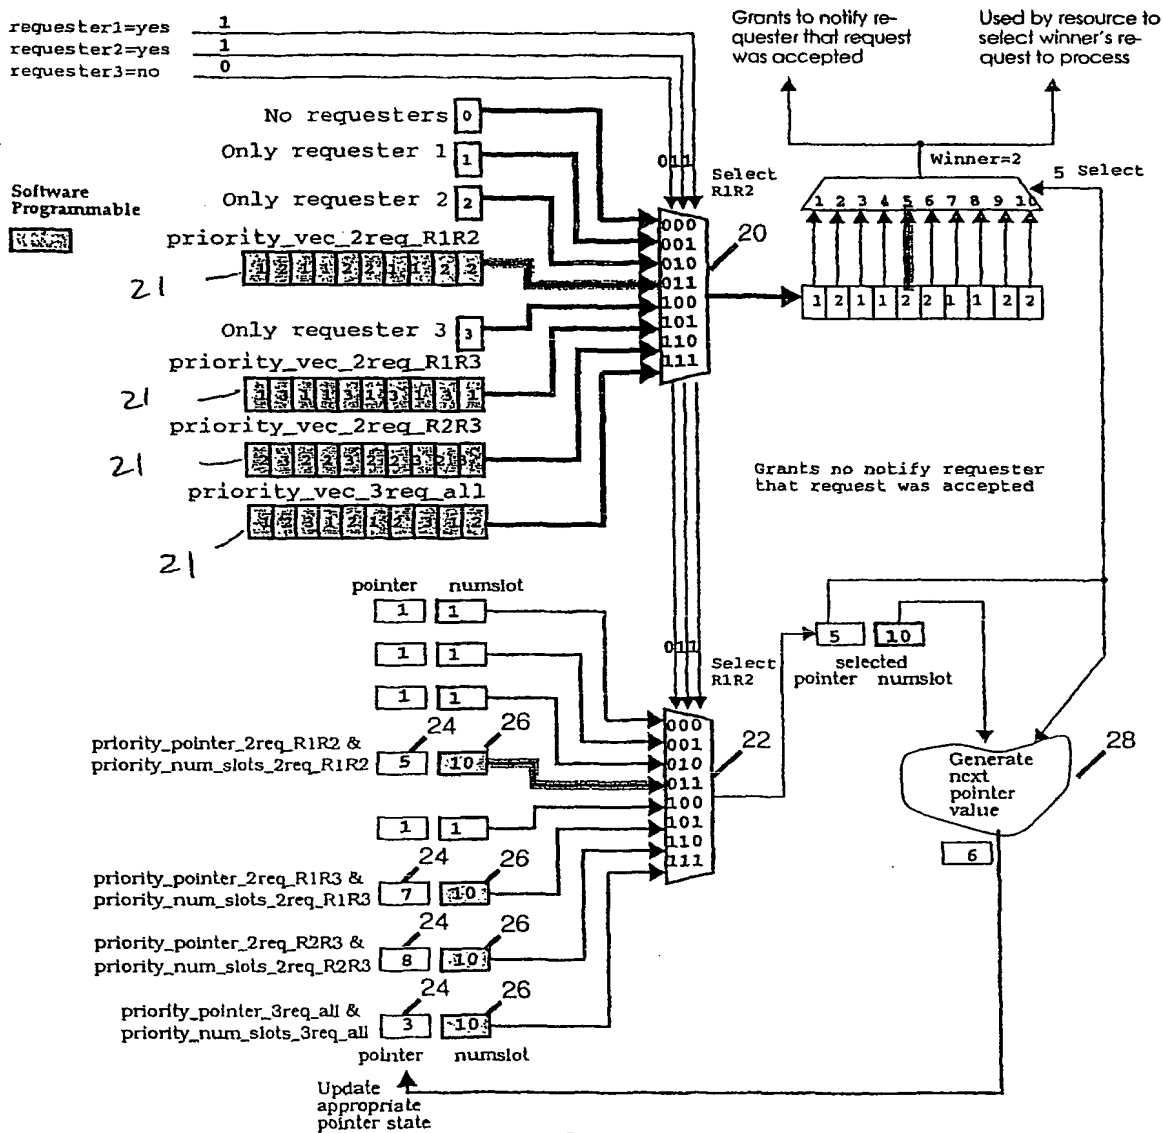


FIG. 2

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
2 May 2002 (02.05.2002)

PCT

(10) International Publication Number
WO 02/035345 A2

(51) International Patent Classification⁷: G06F 9/00

(21) International Application Number: PCT/US01/51236

(22) International Filing Date: 25 October 2001 (25.10.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/697,222 26 October 2000 (26.10.2000) US

CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, M/S PALO1-521, Palo Alto, CA 94303 (US).

Published:

— without international search report and to be republished upon receipt of that report

(72) Inventors: NGUYEN, Hien, H.; 282 Grove Street, Apt. 1, Auburndale, MA 02466 (US). MORRIER, Donald, M.; PMB 312, 119 Drum Hill Road, Chelmsford, MA 01824 (US).

(48) Date of publication of this corrected version:
6 November 2003

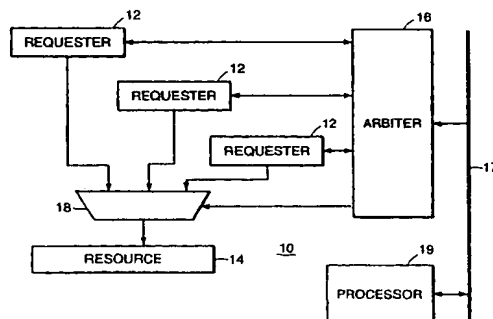
(74) Agents: ASHER, Robert, M. et al.; Bromberg & Sunstein LLP, 125 Summer Street, Boston, MA 02110-1618 (US).

(15) Information about Correction:
see PCT Gazette No. 45/2003 of 6 November 2003, Section II

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: ARBITRATION METHOD AND APPARATUS



(57) Abstract: An arbitration method and mechanism assigning priority to one of a combination of requesters. A priority vector for the combination of requesters has a sequential list of requester identifiers. Upon receiving requests from the combination of requesters, the corresponding priority vector is referred to and the next requester identifier on this list is selected. Priority is awarded to the requester corresponding to that requester identifier. A priority vector is preferably provided for each possible combination of multiple requesters. The list of requester identifiers in each priority vector may be advantageously programmable. Preferably, all requester identifiers for a single requester are not provided in a contiguous clump within a given priority vector.

WO 02/035345 A2

ARBITRATION METHOD AND APPARATUS

Background of the Invention

5 The present invention relates to an arbitration mechanism for awarding priority to one of a plurality of requesters, in particular, such a mechanism which guarantees a minimum software programmable slice of available bandwidth to each requester.

Where there is a single resource in the data processing system whether that resource be a bus, an SRAM, a FIFO, etc., an arbitrator is required to control access to
10 that resource. In accordance with many arbitration schemes, each requester has an assigned priority. Requesters with the same priority are given access in an equitable or round-robin fashion. Typically, requests with higher priority are given preferred access, preempting requests with lower priority. Additional mechanisms have been developed, that when added on top of such a system, prevent requesters with low priority from being
15 starved from access to the resource.

Summary of the Invention

In accordance with an arbitration method embodiment of the invention, requests
20 for access to a resource are received from a combination of requesters. A priority vector associated with that combination of requesters is referred to in order to identify the requester to be awarded priority. Requester identifiers each corresponding to one of the requesters are listed sequentially in the priority vector. The number of requester identifiers in the vector exceed the number of requesters in the combination. A pointer
25 identifies a next requester identifier on the list. The next requester identifier is selected and the requester corresponding to that requester identifier is awarded priority. The pointer may be incremented sequentially through the list of requester identifiers. Programmability can be made available for the total number of requester identifiers in the priority vector and for the number of requester identifiers corresponding to each
30 requester.

According to a further embodiment of the arbitration method, a plurality of priority vectors are provided, one for each possible combination of multiple requesters that may seek access to the resource. The priority vector is a sequence of requester

identifiers, each requester identifier corresponding to one of the requesters. When requests for access are received from the combination of requesters, the priority vector corresponding to the given combination of requesters is used for selection of the next requester identifier. Priority is awarded to the requester corresponding to the next
5 requester identifier.

In accordance with embodiments of an arbitration mechanism of the invention, a selector, responsive to requests for access from the requesters in a given combination of requesters, ascertains a next requester identifier from the sequence of requester identifiers in the priority vector corresponding to the given combination of requesters.

10 Further, a pointer may be included for pointing to one of the requester identifiers in a priority vector. The sequence of requester identifiers may be held in a register.

A data processing system embodiment of the invention includes a resource, a plurality of requesters and an arbiter in communication with the requesters. The arbiter includes a selector responsive to the requests for ascertaining a next requester identifier
15 from an appropriate priority vector and awarding priority to the requester corresponding to a next requester identifier.

Since each requester will have at least one requester identifier in the priority vector used to award priority, even low priority requesters, will be able to gain some minimal access to the resource. The ability to assign a percentage of bandwidth to each
20 requester by adjusting the number of corresponding requester identifiers in priority vectors is an advantage of various embodiments of the invention.

Other objects and advantages of the invention will become apparent during the following description of the presently preferred embodiments of the invention taken in conjunction with the drawings.

25

Brief Description of the Drawings

FIG. 1 is a block diagram of a data processing system of an embodiment of the present invention.

FIG. 2 is a schematic block diagram illustrating an arbitration method and

mechanism of embodiments of the present invention.

Detailed Description of the Preferred Embodiments

Referring now to FIG. 1, in the data processing system 10 it is often the case that a plurality of requesters 12 may at the same time request access to a single resource 14. The resource may be any of a variety of components such as a bus, an SRAM, a FIFO, a network communication port, a graphic processing pipeline, DRAM memory banks, a data structure in a data base program, etc. In order to decide which requester receives access to the single resource 14, an arbiter 16 is provided.

Arbiter 16 receives requests for access from any of a number of requesters. The arbiter 16 awards priority to one of the requesters. In a preferred embodiment, arbiter 16 sends a select signal to multiplexor 18 which selects the data from the requester awarded priority. The selected data proceeds to the resource 14. Also in this embodiment, arbiter 16 informs the requester awarded priority that its request was granted. Preferably, the arbiter is programmable. A processor 19 is shown for configuring the arbiter through a bus 17.

A preferred embodiment of an arbiter of the present invention is illustrated in FIG. 2. The arbitration method is required when more than one requester 12 seeks access to a single resource at one time. For any given combination of requesters, the arbiter 16 has a priority vector 21 with a sequential list of requester identifiers. A requester identifier corresponds to one of the requesters. In a preferred embodiment, the number of requester identifiers in the list of the priority vector 21 exceeds the number of requesters in the given combination. In accordance with further embodiments, each possible combination of multiple requesters is associated with a priority vector 21. For each priority vector 21, a pointer 24 or equivalent designator identifies the next requester identifier on the list. For the given combination of multiple requesters, the associated priority vector is referred to and the next requester identifier from that priority vector is selected. Priority is awarded to the requester corresponding to the next requester identifier. The pointer 24 is then changed to point to the next requester identifier along

the sequence of requester identifiers. As described, the arbitration method may be performed entirely in software or in a combination of software and hardware as described below with respect to FIG. 2.

The arbitration mechanism includes a selector responsive to requests for access.

- 5 The selector is responsible for ascertaining the next requester identifier in the appropriate priority vector for awarding priority to the requester corresponding to the next requester identifier. The selector in FIG. 2 includes a multiplexor 20 for selecting the appropriate priority vector and a multiplexor 22 for selecting the appropriate pointer. The multiplexor 20 receives an input for each possible requester indicating whether a request
- 10 is pending for the single resource. The embodiment shown is for a system with a maximum of three requesters. Therefore, there is a priority vector 21 for all three requesters and three additional priority vectors 21 for each possible combination of two requesters. Based on which requesters have requests pending, the multiplexor 20 immediately selects the corresponding priority vector 21. In the example of the
- 15 illustration, the first and second requesters have pending requests and the priority vector for requester 1 and requester 2 is selected.

- Priority is awarded to the requester corresponding to the next requester identifier in the priority vector. The priority vector is a sequential list of requester identifiers; each requester identifier corresponding to one of the requesters. In the illustration, the
- 20 requester identifier is simply the number 1, 2 or 3 corresponding respectively to requester 1, requester 2 or requester 3. In the priority vector for requester 1 and requester 2, the list includes requester identifiers corresponding to the two requesters 1 and 2. In the illustration, the priority vector for requester 1 and 2 split the priority 50/50 between requester 1 and requester 2, each one having five requester identifiers in the list. In the
- 25 case where a requester has more than one requester identifier in the list, it is desirable not to clump all of its requester identifiers together as such arrangement may inordinately starve the other requester(s) out over a clump of time. Instead, it is preferred for each requester having more than one requester identifier, that its requester identifiers not be contiguous all to one another. This distributes the grant of priority to the various

requesters more evenly over time.

The next requester identifier in a priority vector is indicated by a pointer 24. The multiplexor 22 receives the inputs indicating which requesters have a pending request and uses those inputs to immediately select the pointer corresponding to the selected
5 priority vector. The pointer 24 is used to ascertain the next requester identifier which in turn determines which requester is awarded priority. The award of priority is signaled back to the requester so it knows its request is granted. The award of priority is also used by the resource to select the winning request and process it. In FIG. 1, this is through multiplexor 18. Upon using a pointer 24 to select the award of priority, the
10 pointer needs to be shifted to point to the following requester identifier on the priority vector. This may be accomplished by simply incrementing the pointer. A vector length indicator (numslot) 26 accompanies the pointer for each priority vector. When the total number of requester identifiers in the priority vector have been traversed by the pointer as determined by the vector length indicator, the pointer will wrap around to the first
15 requester identifier on the list. A pointer control module 28 receives the pointer and associated vector length indicator 26 so as to generate the next pointer value and substitute it into the pointer 24. The pointer control module 28 can be formed in hardware, software or anything in between.

In accordance with an embodiment of the present invention, the number of
20 requester identifiers in the priority vector exceeds the number of requesters. Thus, for a system with three possible requesters, the priority vector for all three requesters exceeds three requester identifiers. The more requester identifiers in a priority vector, the smaller the minimum bandwidth which can be assigned to a low priority requester. A priority vector with 10 requester identifiers can assign as little as one requester identifier to a
25 requester giving it approximately a 10% bandwidth. A priority vector with 16 requester identifiers permits assigning a bandwidth as small as about 6% to a low priority requester.

In accordance with an embodiment of the invention, the number of requester identifiers in any given priority vector is programmable. Software instructions can be

used to set the vector length for each of the priority vectors. Further programmability is also available in that the proportion of requester identifiers assigned to the requesters can be varied for each priority vector. Although each priority vector can be individually programmed with a desired sequence and proportion of requester identifiers, it is

- 5 generally preferable to determine a desired percentage bandwidth for each requester and have the various priority vectors attempt to achieve that goal. A system that desires to provide about 40% access to requester 1, about 40% access to requester 2 and about 20% access to requester 3 with a vector length of 10 for each priority vector is shown in FIG.
2. The priority vector for all three provides four requester identifiers for requester 1, 10 four requester identifiers for requester 2 and two requester identifiers for requester 3. The priority vector for requesters 1 and 2 correspondingly provide five requester identifiers for requester 1 and five requester identifiers for requester 2. The priority vector for requesters 1 and 3 provides six requester identifiers for requester 1 and four requester identifiers for requester 3. Similarly the priority vector for requesters 2 and 3 15 provide six requester identifiers for requester 2 and four requester identifiers for requester 3. One might also have chosen to divide these vectors seven to three rather than six to four. As there are a small finite number of requester identifiers, bandwidth is only proportioned approximately. Increasing vector length can make the proportions more precise. Although it is preferred to have consistency of proportions among the 20 priority vectors, it is within the scope of the invention that each priority vector be independently programmed with its own list of requester identifiers without regard to overall proportions. Each requester can thus be given a desired number of requester identifiers in the priority vectors. In each priority vector, the requester identifiers for any given requester are preferably spread across the vector and do not appear in a single 25 contiguous clump.

The arbiter 16 of a presently preferred embodiment can be implemented on an ASIC. The priority vectors may be stored in a register. Each register may be divided into fields. One field provides slots that are filled by the requester identifiers to form priority vectors. A second field indicates the vector length, i.e., the number of slots that

are filled with requester identifiers. A large register can hold several priority vectors in separate fields and have associated fields which indicate the vector lengths for the corresponding priority vectors. The contents of the registers are programmable and may be filled to satisfy the requirements of a user through a processor 19 for example.

- 5 Although shown and described with respect to hardware such as the multiplexors and registers, embodiments of the invention may also be made implementing aspects of the invention in software.

Although the embodiments hereinbefore described are preferred, many modifications and refinements which do not depart from the true spirit and scope of the invention may be conceived by those skilled in the art. For example, the requester identifiers need not be numbers that match their corresponding requester but may be assigned via any desired coding. Shifting the pointer along a priority vector may involve any of incrementing, decrementing or other algorithm for moving the pointer in sequence through a list of requester identifiers. It is intended that all such modifications, including

10 but not limited to those set forth above, be covered by the following claims.

15

WE CLAIM:

1. A method of arbitrating requests for access to a resource for use in a data processing system having a plurality of requesters and the resource, said method comprising:

- 5 receiving requests for access from a combination of said requesters;
referring to a priority vector associated with said combination of requesters, each priority vector having a sequential list of requester identifiers, the number of requester identifiers in the list exceeding the number of requesters in said combination, and a pointer to a next requester identifier on the list;
- 10 selecting the next requester identifier from the priority vector associated with said combination of requesters; and
awarding priority to the requester corresponding to the next requester identifier.

2. The method of claim 1 further comprising setting the number of requester
15 identifiers in the priority vector.

3. The method of claim 2 further comprising setting how many of the requester identifiers correspond to each requester in the combination of requesters associated with the priority vector.

20 4. The method of claim 3 wherein for those requesters having a plurality of requester identifiers in the priority vector, the corresponding requester identifiers for each of those requesters not all being contiguous in the sequential list.

25 5. The method of claim 1 wherein said act of referring comprises identifying from among a plurality of priority vectors, the priority vector associated with said combination of requesters.

6. The method of claim 1 comprising sequentially changing the pointer along the sequence of requester identifiers each time a next requester identifier is selected.

7. A method of arbitrating requests for access to a resource for use in a data processing system having a plurality of requesters and the resource, said method comprising:

providing a plurality of priority vectors, one priority vector for each possible combination of multiple requesters seeking access to the resource, each vector being a sequence of requester identifiers, each requester identifier corresponding to one of the requesters;

receiving requests for access to the resource from a given combination of said requesters;

selecting a next requester identifier from the priority vector corresponding to the given combination of requesters;

awarding priority to the requester corresponding to the next requester identifier.

8. The method of claim 7 further comprising assigning each requester in the combination of requesters corresponding to one of the priority vectors a chosen number of requester identifiers in the priority vector.

9. The method of claim 8 wherein for those requesters having a chosen number greater than one for one of said priority vectors, the corresponding requester identifiers for each of those requesters not all being contiguous in the sequence of requester identifiers.

10. The method of claim 7 further comprising providing a pointer for each priority vector that points to the next requester identifier in the priority vector and sequentially changes along the sequence of requester identifiers each time a next requester identifier is selected.

11. A method for providing at least minimal access to a resource for each requester among a plurality of requesters, said method comprising:

providing a plurality of priority vectors, each priority vector corresponding to a possible combination of said requesters and being capable of holding a set number of requester identifiers in a sequence;

programming each priority vector such that it contains a desired number of requester identifiers corresponding to each of the requesters in the corresponding combination of requesters;

10 pointing to a next requester identifier in each of said priority vectors;

awarding priority to a requester from a combination of said requesters by identifying the next requester in the priority vector corresponding to the combination of requesters and pointing to a requester identifier following the next requester identifier in the sequence of said priority vector.

15

12. The method of claim 11 wherein for each requester having a number of requester identifiers greater than one in one of said priority vectors, the corresponding requester identifiers not all being contiguous in the sequence of requester identifiers of said one of said priority vectors.

20

13. The method of claim 11 further comprising programming the set number of requester identifiers for each priority vector.

14. The method of claim 11 wherein pointing to a requester identifier following the next requester identifier is performed in a wrap around manner upon reaching the end of the sequence.

25

15. An arbitration mechanism for arbitrating requests for access to a resource comprising:

a plurality of priority vectors, each providing a sequence of requester identifiers corresponding to requesters, the requester identifiers in each priority vector corresponding to requesters forming a different combination of requesters;

a selector, responsive to requests for access from the requesters in a given

- 5 combination of requesters, for ascertaining a next requester identifier in sequence in the priority vector corresponding to the given combination of requesters and awarding priority to the requester corresponding to the next requester identifier.

16. The arbitration mechanism of claim 15 wherein each priority vector has a number
10 of requester identifiers exceeding the number of requesters in the respective combination of requesters.

17. The arbitration mechanism of claim 15 wherein for each requester having a
number of requester identifiers greater than one in one of said priority vectors, the
15 corresponding requester identifiers not all being contiguous in the sequence of requester identifiers of said one of said priority vectors.

18. The arbitration mechanism of claim 15 further comprising a processor that sets
the total number of requester identifiers in each priority vector and sets the number of
20 requester identifiers corresponding to any given requester in each priority vector.

19. An arbitration mechanism comprising:
a register for holding a sequence of requester identifiers, said register including
requester identifiers corresponding to each requestor of a combination of requesters,
25 wherein the number of requester identifiers in said register exceeds the number of
requesters in said combination;
a pointer for pointing to one of the requester identifiers in said register; and
a selector, responsive to requests from the requesters in the combination of
requesters, for awarding priority to the requester corresponding to the requester identifier

being pointed at by said pointer; and

means for shifting said pointer along the sequence of requester identifiers.

20. The arbitration mechanism of claim 19 wherein said pointer shifts sequentially
5 through said sequence of requester identifiers in a wrap around manner.

21. The arbitration mechanism of claim 19 further comprising a plurality of said registers, each including requester identifiers corresponding to a different combination of requesters.

10

22. The arbitration mechanism of claim 21 wherein said selector responds to requests from a given combination of requesters by awarding priority using the register having requester identifiers corresponding to the given combination of requesters.

15 23. A data processing system comprising:

a resource;

a plurality of requesters;

an arbiter in communication with said plurality of requesters, said arbiter

including:

20

a plurality of priority vectors, each providing a sequence of requester identifiers, each requester identifier corresponding to one of said requesters, the requester identifiers in each priority vector corresponding to requesters from a different combination of requesters;

25

a selector, responsive to requests for access from the requesters in a given combination of said requesters, for ascertaining a next requester identifier in sequence in the priority vector corresponding to the given combination of requesters and awarding priority to the requester corresponding to the next requester identifier.

30 24. The data processing system of claim 23 wherein each priority vector has a

number of requester identifiers exceeding the number of requesters in the respective combination of requesters.

25. The data processing system of claim 23 wherein for each requester having a
5 number of requester identifiers greater than one in one of said priority vectors, the corresponding requester identifiers not all being contiguous in the sequence of requester identifiers of said one of said priority vectors.

26. The data processing system of claim 23 further comprising a processor that sets
10 the total number of requester identifiers in each priority vector and sets the number of requester identifiers corresponding to any given requester in each priority vector.

02442/00102 178719.1

1/3




FIG. 1

2/3




FIG. 2A

FIG. 2B

3/3




FIG. 2B

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
2 May 2002 (02.05.2002)

PCT

(10) International Publication Number
WO 2002/035345 A3

(51) International Patent Classification⁷: **G06F 9/46**

(21) International Application Number:
PCT/US2001/051236

(22) International Filing Date: 25 October 2001 (25.10.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/697,222 26 October 2000 (26.10.2000) US

(71) Applicant: **SUN MICROSYSTEMS, INC.** [US/US]; 901
San Antonio Road, M/S PALO1-521, Palo Alto, CA 94303
(US).

(72) Inventors: **NGUYEN, Hien, H.**; 282 Grove Street, Apt. 1,
Auburndale, MA 02466 (US). **MORRIER, Donald, M.**;
PMB 312, 119 Drum Hill Road, Chelmsford, MA 01824
(US).

(74) Agents: **ASHER, Robert, M.** et al.; Bromberg & Sunstein
LLP, 125 Summer Street, Boston, MA 02110-1618 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM,
HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK,
LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX,
MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK,
SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian
patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European
patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,
IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF,
CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD,
TG).

Published:


— with international search report

(88) Date of publication of the international search report:
4 March 2004

(15) Information about Correction:

[Continued on next page]

(54) Title: **ARBITRATION METHOD AND APPARATUS**



(57) Abstract: An arbitration method and mechanism assigning priority to one of a combination of requesters. A priority vector for the combination of requester's (12) has a sequential list of requester identifiers. Upon receiving requests from the combination of requesters, the corresponding priority vector is referred to and the next requester identifier on this list is selected. Priority is awarded to the requester corresponding to that requester identifier. A priority vector is preferably provided for each possible combination of multiple requesters. The list of requester identifiers in each priority vector may be advantageously programmable. Preferably, all requester identifiers for a single requester are not provided in a contiguous clump within a given priority vector. Preferably, the number of requester identifiers in the vector exceeds the number of requesters.

WO 2002/035345 A3



Previous Correction:

see PCT Gazette No. 45/2003 of 6 November 2003, Section II

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERNATIONAL SEARCH REPORT

Intern: pplication No

PCT/US 01/51236

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 991 873 A (SETO KOSHI ET AL) 23 November 1999 (1999-11-23) abstract	1-6, 11-26
A	column 1, line 10 -column 2, line 23 column 2, line 66 -column 6, line 26 figures 1A,1B,2,3A	7-10
X	EP 0 658 841 A (IBM) 21 June 1995 (1995-06-21) abstract	1-6, 11-26
A	column 3, line 38 -column 5, line 39 column 6, line 19 -column 9, line 20 --- -/--	7-10

☒ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *Z* document member of the same patent family

Date of the actual completion of the international search

29 August 2003

Date of mailing of the international search report

05/09/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Kesting, V

INTERNATIONAL SEARCH REPORT

Internationa
plication No
PCT/US 01/51236

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 443 782 A (AMERICAN TELEPHONE & TELEGRAPH) 28 August 1991 (1991-08-28) abstract	7,11,13, 14
A	column 1, line 47 -column 2, line 28 column 4, line 42 -column 10, line 37 column 14, line 46 - line 52 column 16, line 46 -column 17, line 19 column 20, line 17 - line 45 -----	1,15,19, 23

INTERNATIONAL SEARCH REPORT

Information on patent family members

Internal Application No
PCT/US 01/51236

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5991873	A	23-11-1999	JP 3144842 B2	12-03-2001
			JP 5046413 A	26-02-1993
			KR 9603045 B1	04-03-1996
EP 0658841	A	21-06-1995	US 5487170 A	23-01-1996
			EP 0658841 A2	21-06-1995
			JP 7200318 A	04-08-1995
EP 0443782	A	28-08-1991	US 5272697 A	21-12-1993
			CA 2036603 A1	24-08-1991
			CA 2036603 C	21-11-1995
			EP 0443782 A2	28-08-1991
			JP 2527851 B2	28-08-1996
			JP 5304534 A	16-11-1993

THIS PAGE BLANK (USPTO)